

Generating Bijections between HOAS and the Natural Numbers*

John Tang Boyland

College of Engineering and Applied Science
University of Wisconsin–Milwaukee
Milwaukee, Wisconsin USA
boyland@cs.uwm.edu

A provably correct bijection between higher-order abstract syntax (HOAS) and the natural numbers enables one to define a “not equals” relationship between terms and also to have an adequate encoding of sets of terms, and maps from one term family to another. Sets and maps are useful in many situations and are preferably provided in a library of some sort. I have released a map and set library for use with Twelf which can be used with any type for which a bijection to the natural numbers exists.

Since creating such bijections is tedious and error-prone, I have created a “bijection generator” that generates such bijections automatically together with proofs of correctness, all in the context of Twelf.

1 Introduction

Higher-order abstract syntax (HOAS) [9] uses the functions of the meta-logic to represent functions (and related constructors, such as “let”). For example (in each case, I define the canonical identity function `id`):

Traditional Abstract Syntax:

```
t : type.  
var : name -> t.  
lam : name -> t -> t.  
app : t -> t -> t.
```

```
%abbrev id : t = lam x (var x).
```

Higher-Order Abstract Syntax:

```
t : type.  
  
lam : (t -> t) -> t.  
app : t -> t -> t.
```

```
%abbrev id : t = lam ([x] x).
```

Not only does the traditional syntax need a type for “names” (where “x” is a typical instance) but then also must handle the fact that a variable may be undeclared. Furthermore, there are the problems of accidental name clashes and alpha-equivalence: the two functions `lam X (var X)` and `lam Y (var Y)` are different if the names `X` and `Y` are different. The encoding of functions using names has both “junk” and “duplicates.”

Higher-order syntax maintains alpha-equivalence directly but only makes sense in a logic in which the function cannot perform case analysis on its argument. One does not want the abstract syntax of a function to depend on the semantic value of a parameter at run-time!

*Work supported in part by the National Science Foundation (CCF-0702635). The opinions expressed here are not necessarily those of the National Science Foundation or the US Government.

A technique that avoids duplicates is to use nameless terms (de Bruijn terms). In order to avoid junk as well, one uses an “index” on the types [1, 5]:

```

var : nat -> type.
1 : var (s N).
1+ : var N -> var (s N).

term : nat -> type.
var : var N -> term N.
lam : term (s N) -> term N.
app : term N -> term N -> term N.

%abbrev id : term z = lam (var 1).

```

Here `var N` is the type of all legal variables inside N lambda abstractions; it has N elements. When N is zero (written z), this type is empty. Similarly `term N` is the type of terms inside N lambda abstractions.

When one reasons with ASTs, it will sometimes be convenient to determine whether two trees are equal or unequal, such as when one forms sets of ASTs for analysis purposes. It is fairly straightforward (if tedious) to define inequality for traditional or nameless terms, but with HOAS, this task is made much more complex because one has to reason about inequality of functions. Furthermore, as I describe in the following section, it is useful to have a bijective mapping from ASTs to the natural numbers. The definition of inequality is trivial given a bijection.

This paper shows how such mappings can be defined in general, and proved correct. I describe a tool I have implemented that produces such mappings and proofs automatically.

In the following section, I review the concept of “adequacy” and show how the desire for adequate encodings of sets and maps motivates the definition of bijections between arbitrary term languages and the natural numbers. Section 3 shows how such mappings can be defined, with particular attention to HOAS and indexed types. Then Section 4 explains how proofs can be constructed to prove the correctness of the bijection. Section 5 describes the implementation of a tool that generates the mapping and correctness proofs in Twelf.

2 An “Adequate” Encoding of (Finite) Sets

Adequacy [3, 4] is an important concept in a logical system—it means that the source concepts (mathematical in nature) are faithfully represented in the logical system. In particular, an adequate encoding is a bijective mapping from the source concept to a target type for the logical representation. That is, the mapping is

total Every instance of the source concept has an encoding.

unique The mapping is a function (“no confusion”).

onto Every instance of target type represents a valid instance of the source concept (“no junk”).

one-to-one Two distinguishable instances of source concept always map to distinguishable instances of the target type (“no loss”).

In this section, I demonstrate the concept of adequacy by showing how it applies when encoding finite sets of natural numbers.

The first and last requirements (“total” and “one-to-one”) are obviously needed for correctness. If a set cannot be represented, or if two different sets are represented the same way, the representation obviously is not faithful.

	$\{0, 5\} \Rightarrow \langle 0, 4 \rangle$	$\{0, 2, 4\} \Rightarrow \langle 0, 1, 1 \rangle$
$\{\} \Rightarrow \langle \rangle$	$\{1, 5\} \Rightarrow \langle 1, 3 \rangle$	
$\{0\} \Rightarrow \langle 0 \rangle$	$\{2, 5\} \Rightarrow \langle 2, 2 \rangle$	$\{2, 3, 4\} \Rightarrow \langle 2, 0, 0 \rangle$
$\{3\} \Rightarrow \langle 3 \rangle$	$\{4, 5\} \Rightarrow \langle 4, 0 \rangle$	$\{4, 11, 96\} \Rightarrow \langle 4, 6, 84 \rangle$

Figure 1: Illustration of an adequate encoding of finite sets of natural numbers.

The middle requirements (“unique” and “onto”) may also seem obvious but are frequently violated in practice. For example, the common practice of representing sets of natural numbers with lists of natural numbers without duplicates violates both requirements: the encoding is not unique because reorderings of a list generates duplicate representations of the same set; neither is it onto because lists of duplicates are “junk” that do not validly represent any set.

What is the cost of having an encoding that is not unique and/or not onto? One must define auxiliary relations. In the case of a lack of uniqueness, one must define an equivalence relation over the representation type. Furthermore, one must prove numerous lemmas that say that equivalence is preserved by all operations on sets (union, intersection, inclusion, containment, etc). Since equivalence is non-trivial when the encoding is not unique, these lemmas are also not trivial. Furthermore, equivalence “invades” all the uses of sets in an application (i.e., a logical system that relies on sets): every theorem concerning application relations touching on sets will need to incorporate equivalence.

In the case of an encoding with junk, one must define a validity relation, true for the subset of instances of the representation type that have a source representation. Such a validity relation is the analog of a data structure invariant. Again, it will be necessary to prove that validity is preserved by set operations and by all application relations using sets. If a set is in an “input” argument, there must be an additional “validity” requirement on input. Similarly, an “output” set value must be accompanied by a check of validity. If an encoding is neither unique nor onto, one will need both equivalence and validity.

Therefore, in order not to burden the application with equivalence and/or validity, I defined an encoding of sets of natural numbers that is fully adequate.¹ A set is represented by a sequence of numbers, the first of which is the smallest element in the set and the remainder are the counts of *missing* numbers between adjacent (in sorted order) numbers in the set. Examples of this encoding are shown in Figure 1.

This technique for forming sets does not immediately apply to other types: how would one form (say) a set of strings, of sets, or of HOAS terms? The difficulty is that one needs to be able to count the elements “between” any two elements in the domain. This requirement is satisfied if one has a bijection from the element type to the natural numbers. Using the bijection, one can easily form sets of the element type by mapping back and forth to integers and using the existing “set” signature.

Sets of course can be seen as a special case of maps, and indeed my own set signature is created by specializing a map signature for the “unit” result type. Thus a bijection between a type and the natural numbers is useful for defining adequate encodings of (finite) mappings as well. This then is my motivation for exploring bijections between arbitrary term languages and the natural numbers.

¹I expect others have done so as well, but there is apparently no accepted library of such signatures for Twelf. My signature is released to the public domain. This release includes other useful signatures, such as an adequate encoding of positive rational numbers.

```

nat : type.                                %abbrev 0 : nat = z.
z : nat.                                  %abbrev 1 : nat = s 0.
s : nat -> nat.                            ...

plus : nat -> nat -> nat -> type.          %abbrev 0+0=0: plus 0 0 0 = plus/z.
plus/z : plus z Y Y.                     %abbrev 1+0=1: plus 1 0 1 = plus/s 0+0=0.
plus/s : plus (s X) Y (s Z)              %abbrev 0+1=1: plus 0 1 1 = plus/z.
      <- plus X Y Z.                      ...

times : nat -> nat -> nat -> type.         %abbrev 0*0=0: times 0 0 0 = times/z.
times/z : times z X z.                   ...
times/s : times (s X) Y Z
      <- plus T Y Z

mingle : nat -> nat -> nat -> type.       %abbrev 0$0=0: mingle 0 0 0 = ... .
...                                       ...

```

Figure 2: Natural numbers, some operations and some abbreviations.

3 Mapping Terms to Natural Numbers

One of the delightful results of set theory is that the cardinality of the set of pairs of natural numbers is the same as the cardinality of the natural numbers themselves, in other words, that the set of pairs of natural numbers can be arranged in a sequence with a particular starting point. A number of different simple mappings can be defined. I use the INTERCAL [11] “mingle” operation (written as infix “\$”), which interleaves the bits of two numbers to create a (unique) result:

0\$0 = 0	0\$1 = 1	0\$2 = 4	0\$3 = 5
1\$0 = 2	1\$1 = 3	1\$2 = 6	1\$3 = 7
2\$0 = 8	2\$1 = 9	2\$2 = 12	2\$3 = 13
3\$0 = 10	3\$1 = 11	3\$2 = 14	3\$3 = 15

Figure 2 shows selected parts of the “nat” signature used in this paper as well as some abbreviations used in the generator. The full definition of mingle is not shown.

3.1 Simple Mappings

A finite term type does not have a bijection to the natural numbers (of course) but rather only to a subset of the same size. finite types (such as booleans) are mapped to consecutive natural numbers $0..n-1$.

When a term type has multiple constructors, then the mapping I define distinguishes the “finite” constructors (whose instances are enumerated) from the infinite ones. For example:

```

natlist : type.

natlist/0 : natlist.                      L(natlist/0) = 0
natlist/+ : nat -> natlist -> natlist.    L(natlist/+ n l) = 1+n$L(l)

```

When there are multiple “infinite” constructors, the mapping results are distinguished through multiplication, as seen in the following mapping for my rational type (based on continued fractions):

rat : type.	Bijection:	Not a bijection:
whole : nat -> rat.	$R(\text{whole } n) = 2n + 0$	$R'(\text{whole } n) = 2n + 1$
frac : nat -> rat -> rat.	$R(\text{frac } n \ r) = 2(n\$R(r)) + 1$	$R'(\text{frac } n \ r) = 2(n\$R'(r)) + 0$

Generating such mappings is fairly easy, although even with such simple types, one must be careful: the superficially similar mapping R' is not onto. This fact can be seen since a t for which $R'(t) = 0$ would need to be of the form $t = \text{frac } z \ r$, where $R'(r)$ would again have to be zero.

Ignoring the latter issue, it is simply then a matter of handling data type definitions with a mixture of finite and infinite constructors, each of which may use subterms of finite type, and whose infinite constructors may have any number of subterms of any infinite type. The situation is more interesting in the presence of HOAS.

3.2 Mapping functions

With higher-order abstract syntax (HOAS), the mapping essentially connects each variable with a small integer (in the range $0..N - 1$ where N is the number of variables in scope). For example, consider the pure lambda calculus defined previously:

t : type.	$M_N(\text{lam } \lambda x. t) = N + 2(M_{N+1}(t)) + 0$
	where $M_X(x) = N$
lam : (t -> t) -> t.	$M_N(\text{app } t_1 \ t_2) = N + 2(M_N(t_1) \$ M_N(t_2)) + 1$
app : t -> t -> t.	

The extra parameter N on the mapping indicates the number of free variables in scope. One uses $M_0(t)$ to determine the mapping for a fully-bound term t . Here the definition of the mapping uses “where” clauses to express “hypothetical judgments.” There is no “var” case since there is no “var” constructor for the type. The alert reader will notice that the mapping of a variable does not change when it is used inside nested lambda abstractions; it uses what Pierce [10] calls “de Bruijn” levels, not “de Bruijn” indices.

The nice thing about the mapping is that it preserves alpha-equivalence—two terms at the same level map to the same integer if and only if they are the same (under alpha equivalence). A consequence of the bijection is that it is easy to define “not equal” over higher-order terms via the mapping; such a relation is much harder to define directly.

Here the example uses variables of an infinite type; one can also define higher-order terms over finite types, but I have not seen an application of this. Therefore because it complicates the framework (the cardinality of the finite type changes depending on the context), my implementation does not handle variables of finite type.

3.3 Indexed Types

An indexed type [5] uses one term to distinguish different kinds of a second term. A simple example² that uses variables and indexing is the following first-order functional language with recursion and a single constant:

²This term family is a simplification of a typed predicate system alluded to later.

```
term : nat -> type.
```

```
unit : term z.
```

```
lam : (term z -> term N) -> term (s N).
```

```
app : term (s N) -> term z -> term N.
```

```
rec : (term N -> term N) -> term N.
```

Here `term N` is the type of terms that require N arguments before they can be reduced to (first-order) values. There is one canonical value, `unit`; `lam` takes a function that accepts a value and returns a term requiring N arguments and constructs one that now needs $N + 1$ arguments; `rec` is used to create a recursive function of N arguments.

Now the mapping M_V^n is qualified both by a “vector” (finite map of nats to nats) V and by n , the index of the term. The vector keeps track of how many variables of the given index are currently free. Each index has its *own* bijection to the natural numbers, with differing sets of constructors: `unit` for the $n = 0$ case only and `lam` for the $n > 0$ case only. Thus for the $n = 0$ case, the mapping handles one finite constructor (`unit`) and two infinite constructors (`app` and `rec`), but the mapping for non-zero indices ($n + 1$) does not handle any finite constructors and has three infinite constructors:

$$\begin{aligned}
M_V^0(\text{unit}) &= V(0) \\
M_V^0(\text{app } t_1 \ t_2) &= V(0) + 1 + 2 (M_V^1(t_1) \$ M_V^0(t_2)) + 0 \\
M_V^0(\text{rec } \lambda f. t) &= V(0) + 1 + 2 (M_{V+0}^0(t)) + 1 \quad \text{where } M_X^0(f) = V(0) \\
M_V^{n+1}(\text{lam } \lambda x. t) &= V(n+1) + 3 (M_{V+0}^n(t)) + 0 \quad \text{where } M_X^0(x) = V(0) \\
M_V^{n+1}(\text{app } t_1 \ t_2) &= V(n+1) + 3 (M_V^{n+2}(t_1) \$ M_V^0(t_2)) + 1 \\
M_V^{n+1}(\text{rec } \lambda f. t) &= V(n+1) + 3 (M_{V+(n+1)}^{n+1}(t)) + 2 \quad \text{where } M_X^{n+1}(f) = V(n+1)
\end{aligned}$$

Here $V(n)$ applies the vector to n and $V + n = V[n \mapsto V(n) + 1]$ returns the vector with one more n .

3.4 Binding variables

Originally, I defined mappings using hypothetical judgments as suggested earlier in this section. However, separating the variable information into a vector passed explicitly and hypothetical judgments passed implicitly made the proofs of “one-to-one” (in particular) extremely difficult.³ Thus I changed the way variables were handled: the variables are now passed explicitly in a large sequence that is then split in order to find the location of a variable in the sequence. This approach can be seen as another application of the idea of “explicit contexts” [2].

In the end, I needed *both* approaches: one in which a variable is bound with its “level” in the implicit context (the normal Twelf technique) and one in which the level is implicit in the location of the variable in the explicit context list (so called “nolevel”). Then I generate proofs that a use of a variable with a level can be converted into one without and vice versa.

Some of proofs require that a variable is either in the explicit context without a level or in the implicit context with a level but not both. Other proofs require that variables are distinguished from normal terms. Thus the generator defines three judgments only used hypothetically: `var`, `level` and `nolevel` (see Fig. 3).

³The Twelf wiki shows the extreme lengths I went to for a small indexed higher-order term system.

```

var : term K -> type.
level : term K -> nat -> type.
nolevel : term K -> type.

%block block#var      : some {k}      block {x:term k} {v:var x}.
%block block#level    : some {k} {l} block {x:term k} {v:var x} {vl:level x l}.
%block block#nolevel  : some {k}      block {x:term k} {v:var x} {nl:nolevel x}.

```

Figure 3: Uninhabited types used in hypothetical judgments.

```

list : type.

list/0 : list.
list/+ : {k:nat} {x:term k} (var x) -> list -> list.

```

Figure 4: Definition of variable list type.

```

split : list -> term K -> list -> list -> type.
count : nat -> list -> nat -> type.

split/here : split (list/+ _ X _ L) X list/0 L.
split/there : split L X L1 L2 -> split (list/+ _ _ V L) X (list/+ _ _ V L1) L2.

count/0 : count _ list/0 z.
count/= : count K L N -> count K (list/+ K _ _ L) (s N).
count/!= : count K L N -> nat'ne K' K -> count K (list/+ K' _ _ L) N.

```

Figure 5: Operations on variable lists.

```

map : {K:nat} list -> term K -> nat -> type.

map/level : level X N -> map K _ X N.
map/nolevel : nolevel X -> split H X _ L -> count K L N -> map K H X N.
map/unit : count z H NV -> map z H unit NV.

map/app : map (s z) H A0 N1 -> map z H A1 NA1 -> mingle N1 NA1 N2 ->
  count z H NV -> times 2 N2 TN -> plus 0 TN PN -> plus 1 NV NA ->
  plus PN NA N -> map z H (app A0 A1) N.

map/rec : ({x} {v} nolevel x -> map z (list/+ z x v H) (A0 x) N1) ->
  count z H NV -> times 2 N1 TN -> plus 1 TN PN -> plus 1 NV NA ->
  plus PN NA N -> map z H (rec A0) N.

map/lam : ({x} {v} nolevel x -> map K (list/+ z x v H) (A0 x) N1) ->
  count (s K) H NV -> times 3 N1 TN -> plus 0 TN PN -> plus 0 NV NA ->
  plus PN NA N -> map (s K) H (lam A0) N.

map/app1 : map (s (s K)) H A0 N1 -> map z H A1 NA1 -> mingle N1 NA1 N2 ->
  count (s K) H NV -> times 3 N2 TN -> plus 1 TN PN -> plus 0 NV NA ->
  plus PN NA N -> map (s K) H (app A0 A1) N.

map/rec1 : ({x} {v} nolevel x -> map (s K) (list/+ (s K) x v H) (A0 x) N1) ->
  count (s K) H NV -> times 3 N1 TN -> plus 2 TN PN -> plus 0 NV NA ->
  plus PN NA N -> map (s K) H (rec A0) N.

```

Figure 6: Mapping for the example indexed type.

The two operations on lists are given in Fig. 5. The `split` operation locates a variable within a list, and `count` determines its level, by counting how many variables (of its type) are in the context before it. Here `nat` is the inequality operator for natural numbers.

Finally, Fig. 6 gives the mapping as defined by the generator (cleaned up for readability). The basic structure is the same as seen for the M definition earlier except that it uses variable lists; just as with M , there are separate cases for the zero and non-zero term types. The `map/level` case is used internally in proofs as described in the following section.

4 Proving Correctness

In order for the bijection to be used in a proof system, one needs a proof that the mapping is indeed a bijection. To wit, one needs four proofs—one of each of the four aspects of a bijection outlined in Section 2. This section describes how such a proof can be constructed in Twelf.

4.1 Simple Mappings

Ignoring HOAS and indexed terms, three of the theorems (technically “meta-theorems” in Twelf) are straightforward: totality is proved as a simple “effectiveness” lemma, uniqueness is proved using the uniqueness of the operations (addition, multiplication, mingle) that are used to define it and “one-to-one” is proved using the uniqueness of remainder arithmetic (to distinguish the infinite cases). The last theorem (one-to-one) is long for complex types with many constructors, because one has to check every case against every other case, but each conflicting case is simply a matter of using arithmetic to produce a contradiction. My “nat” signature provides all the necessary theorems on addition, multiplication and division.

With the proof of “onto,” there remains a termination issue: the theorem is demonstrating that given a natural number, one can produce a term that maps to it. This theorem is of course recursive (inductive) in its natural number argument. Thus in a recursive (inductive) call to the theorem, the number must decrease. The result of the mingle operation is greater than its inputs in all cases except when the result is 0 or 1. Thus, the main cases of the theorem may be restricted to only work for $n > 1$, and then to have special cases for $n = 0, 1$. Even these cases may cause unfounded recursion (induction) if the mapping was defined incorrectly (as in the case for R' shown earlier).

The proof of “one-to-one” is the proof that two terms mapping to the same number are the same term. Since equality of terms is defined in terms of identity, the generator needs to define a lemma for each constructor that says that if all the pieces are equal then the result is equal too. Without variables, these lemmas are all trivial to prove. Even with variables, the lemmas are not complex, but if the variable type is indexed, there is an interesting complication (see Sec. 4.3).

4.2 Variables

When proving totality, it makes a difference whether the source is a regular term or a variable (hypothetical in the context). As anyone who has used variables in Twelf knows, one cannot use normal case analysis to distinguish variables. Instead one needs to define an auxiliary predicate that distinguishes them, and then an effectiveness lemma that indicates that the case analysis is always possible. Furthermore, the effectiveness lemma needs to be in its own context.

As mentioned in the previous section, the context is not used to define the mapping for variables; rather an explicit list of variables is used to compute the mapping. This technique however is difficult to handle with the totality and uniqueness theorems: how does one know that every variable is present at least once in the list (totality) and no more than once (uniqueness). Twelf’s “regular worlds” do not give a way to connect a relation’s parameter (the list of variables in this case) and the context. Thus there would seem to be an impasse.

The original technique (in which I used hypothetical judgments to bind each variable to its “level,” the number to which it should be mapped) had none of these problems—totality and uniqueness fall out immediately. This suggested a solution: for totality, the proof would bind the variable in the context to its level, and then after all its uses had been taken care of, remove the level binding and replace it with the technique of using the parameter list. The lemma for replacing a mapping dependent on a level in the context with one that uses the variable list is defined as follows:⁴

⁴ I use the %theorem syntax because (1) it distinguishes meta-theorems from normal relations and (2) it is clearer for non-Twelf experts.

```

%theorem map-remove-term-level:
  forall*
    {K1} {K2}
    {FH:{x:term K1} (var x) -> list}}
    {F:term K1 -> term K2}
    {L:nat} {N:nat}
    {H1:list} {H2:list}
  forall
    {FM:{x:term K1} {v:var x} level x L -> map K2 (FH x v) (F x) N}
    {FS:{x:term K1} {v:var x} split (FH x v) x H1 H2}
    {C:count K1 H2 L}
  exists
    {FM':{x:term K1} {v:var x} nolevel x -> map K2 (FH x v) (F x) N}
  true.

```

Ignoring the implicit (forall*) arguments, the theorem has three inputs (FM, FS, C) and one output (FM'). Thus FM says that the mapping can be computed for the term (F x) that may use x as long as there is a level defined for x. The inputs FS and C say that x is in the parameter list and that the count yields the same value as the level. The result says the mapping does *not* need the level, it can be done merely using the list. This lemma is straightforward to prove using induction: the base case for changing a variable to use the context, and the inductive case for lam are shown:

```

- : map-remove-term-level ([x] [v] [l] map/level l) FS C
    ([x] [v] [n] map/nolevel n (FS x v) C).

- : map-remove-term-level
    ([x] [v] [l] map/lam ([x0] [v0] [n0] F x0 v0 n0 x v l) (FC x v) D3 D4 D5 D6)
    FS C
    ([x] [v] [n] map/lam ([x0] [v0] [n0] G x0 v0 n0 x v n) (FC x v) D3 D4 D5 D6)
    <- ({x0} {v0} {n0} map-remove-term-level
        (F x0 v0 n0)
        ([x] [v] split/there (FS x v)) C
        (G x0 v0 n0)).

```

This lemma (similar to Cray's cut lemma for explicit contexts [2]) enables totality to be proved.

When proving uniqueness, there is already a mapping (indeed, two mappings) using the list ("no level") technique, so the generator uses a lemma that puts the levels back before looking at the bodies of the mappings. For this to work, it is essential to use the fact that the variable occurs in only one place in the list. The generator essentially proves the "reverse" lemma for which FM' along with FS and C are inputs and FM is the output. This is why the type of FS carefully uses H1 and H2 as the type of the suffix and prefix, so that they are known to be independent of the variable, and thus not including it. This permits one to ensure that a variable cannot have two places within the list.

Variables don't add much to complicate the "onto" and "one2one" theorems, except for more special cases for mapping onto 0 and 1.

4.3 Indexed Types

Indexed types complicate the "one-to-one" theorem, in particular the trivial lemmas that say if one constructs a term with equal parts, the results are equal (identical). The problem is that the lemmas cannot

always be typed. Suppose there were another constructor for our type in which an index variable does not occur in the result type:

```
exists : {M:nat} (term M -> term N) -> term N.
```

Notice that M is not used in the result type. (The name `exists` is used because if one defines an indexed predicate syntax with existentials, the existential term will have this form.) Then, the lemma that shows that equal parts produce an equal whole would have the following definition:

```
%theorem exists-respects-eq:
  forall* {A0} {B0} {C}
    {A1:term A0 -> term C} {B1:term B0 -> term C}
  forall
    {EQ0:nat'eq A0 B0}
    {EQ1:{x : term ???} eq (A1 x) (B1 x)}
  exists
    {EQ2:eq (exists A0 A1) (exists B0 B1)}
  true.
```

The problem here is that the type of x must be `term A0` and `term B0` at the same time. Of course, these types are the same (as demonstrated by `EQ0`) but the type system doesn't "know" this fact. This conundrum is solved by defining an equality predicate directly on the functions:

```
func-eq : (term K3 -> term K1) -> (term K4 -> term K2) -> type.
func-eq/ : func-eq F F.
```

This extra layer requires another lemma in the mutual induction of lemmas in proving the result:

```
%theorem map-one2one/func-eq:
  forall*
    {K1:nat} {K'1:nat} {K2:nat} {K'2:nat} {H:list}
    {F1:term K1 -> term K2}
    {F2:term K'1 -> term K'2}
    {N1:nat} {N2:nat}
  forall
    {E1:nat'eq K1 K'1}
    {E2:nat'eq K2 K'2}
    {FM1:{x} {v} nolevel x -> map K2 (list/+ K1 x v H) (F1 x) N1}
    {FM2:{x} {v} nolevel x -> map K'2 (list/+ K'1 x v H) (F2 x) N2}
    {EN:nat'eq N1 N2}
  exists
    {EF:func-eq F1 F2}
  true.
```

Then `exists-respect-eq` is changed to use `func-eq`: `{EQ1: func-eq A1 B1}`.

The biggest problem however with handling indexed types is if they are not "uniform." I call an indexed type *uniform* if all instances of the type (for different indices) have the same cardinality. I stumbled on a practical example of a non-uniform indexed type in typed arguments:

```
argtype : type.
argtype/0 : argtype.
argtype/+ : nat -> argtype -> argtype.
```

```

actuals : argtype -> type.
actuals/0 : actuals argtype/0.
actuals/+ : term K -> actuals A -> actuals (argtype/+ K A).

```

This type lets one ensure that in a call the formal types (not shown) and the actual types match (“correct by construction”). The `actuals` indexed type is however not uniform because the type `actuals A` is sometimes finite (cardinality 1 in fact) and sometimes infinite. The indexed type `var` (defined on page 22) is also not uniform.

In the end, I decided that I would not try to support non-uniform indexed types. If one wishes to use the bijection generator, one must alter the source type, for example by changing the constructor to make it infinite:

```

actuals/0* : nat -> actuals argtype/0.
%abbrev actuals/0 = actuals/0* z.

```

Thus the problem is pushed back to the source type which now has “junk” in it (non-zero numbers in empty argument lists).

The counter-intuitive result is that my generator works on HOAS, but not on the de Bruijn version of the same syntax.

4.4 Complex Type Families

In Twelf, there is no distinction between type families used to form terms, and those used to form relations among terms, and indeed those used to express meta-theorems about relations among terms. What should a bijection generator do when presented with a type such as the following?

```

plus : nat -> nat -> nat -> type.

plus/z : plus z N N.
plus/z : plus X Y Z -> plus (s X) Y (s Z).

```

Given the limitation just described, the answer is easy: “nothing, because `plus` is not uniform.” Furthermore, as described in the following section, my generator permits only a single index, not three indices as seen here.

However, even looking beyond the current limitations of my tool, any generator must be able to determine the cardinality of a type in order to correctly define the mapping. Determining the cardinality of a type is obviously undecidable. Thus at some point the tool must break down, which is unsurprising.

5 Implementation

In this section, I describe the tool that generates the mapping and theorems for a source signature. After outlining how to use it, I describe its limitations and then to what extent it can be thought of as “correct.”

5.1 Running the tool

The generator reads in a signature `required.elf` that defines natural numbers, the uninhabited `void` type, and a boolean type. It assumes the existence of all the relations and theorems from my publicly available `nat.elf` and `natpair.elf` signatures.

It then reads in the argument which should define only the terms for which a bijection is desired. Abbreviations are also permitted, and expanded where they occur. If it notices a problem, it halts execution (or simply throws an exception ...).

It then generates a large signature on standard output: this defines the mapping and the four theorems for each type. Along the way it defines many auxiliary relations and lemmas. To avoid name clashes, it ensures that (except for the abbreviations of the form shown in Fig. 2) all named entities have the hash sign (#) somewhere in the name.

The generated signature for the simple indexed type on page 26 has over 250 definitions in over 60KB of text.

5.2 Power

I wrote the tool in order to generate the bijection for my permission system which has five sorts of terms, three of which are defined together (permission, unit permission and formula) and two which are independent of these (object, wrapping a nat; fraction, wrapping a rational number) but all five of which are defined under a generalized term type `gterm`. This type thus is indexed with a finite index type. Along with this, I define predicate formal and actual types using indexing to ensure parameter matching. Predicate calls are an instance of a formula and hence a `gterm`. On the other hand, predicate formal parameters can be any `gterm`. The predicate system has the equivalent of `rec` (see above) indexed on the argument types. The generated signature has over 1000 definitions in over 300KB of text.

Thus I had a fairly large and complex system to test with. For now, I only need the bijection for this system, but I wanted to be able to add new kinds of permissions and be able to regenerate the mapping and proofs automatically. The system is not much tested beyond this test case and has a number of known limitations.

5.3 Limitations

As explained previously, no tool could do the task for all type families because in general the task is undecidable. But, the tool is fairly limited. In particular:

- Variables can only be of infinite type. That is, any HOAS function must be over an infinite parameter type.
- Indexed types must have only a single index.
- Indexed types must be uniform.
- The index type must not itself be indexed or be defined with HOAS.
- Explicit abstraction $\{X:xxx\} \dots$ must be explicitly typed (`xxx` must be specified) and is only permitted if `xxx` is an index type and `X` does *not* occur in the result type of the constructor (e.g., `exists` on page 31).
- Unbound variables are permitted only if the variable represents an index, the variable is *not* attributed with a type and the variable is used in the result type of the constructor (e.g., `actuals/+` on page 32).
- The result type of an indexed constructor can only use one level of pattern matching if the index is infinite. For example, a constructor for indexed type `term : nat -> type` (the index, `nat` is infinite) may have type $\dots \rightarrow \text{term } (s \ N)$ but not $\dots \rightarrow \text{term } (s \ (s \ N))$ which would require multiple levels of pattern matching. If the index type is finite, no such restriction exists because the tool enumerates all instances of finite index types.

5.4 Correctness, or rather incorrectness

I am reasonably confident that the tool will work well for simple small examples, but if the type family is complex, the tool may reject it, die with a match failure or even get caught in an infinite loop. Furthermore, the generated mappings and theorems may not type check and even if so, the meta-theorems may fail their totality checks.

All these failings are acceptable because the tool does not pretend to be in the trusted core of a system. If the generated mapping and theorems work, then they are correct on their own.

6 Related Work

Gödel numbers are used to represent formal terms as natural numbers so that number theory can be applied to formal logic. For this purpose the mapping need only be total, unique and one-to-one; specific mappings I have seen used in logic are never “onto” and need not be. Variables are usually handled using a fixed (possibly infinite) set of distinct symbols. There is no attempt to maintain alpha-equivalence.

Unlike Twelf, the Coq proof system [6] has an extensive library. It appears that the library includes two implementations of finite sets: one using sorted lists and one using balanced binary trees. Neither implementation is “adequate” in the sense described in Section 2: sorted lists have an extra invariant that is preserved by the operations, and the balanced binary tree technique not only has an invariant but also a non-trivial equivalence relation. I expect that the invariants and probably the equivalence relation are handled through tactics defined for use with the type.

A reviewer observed that the entire generator might be seen as something like a “tactic” in Coq: a rule to help generate proofs. This similarity is intriguing, but one large pragmatic difference is that the tool is intended to be used to generate a signature (Twelf file) that would then be used along with the normal hand-written signatures. In Coq, the tactics are named in the proof scripts and are checked whenever the proof needs to be checked. My generator would have to be much more robust and transparent were it to be used in this way.

As a system that generates proofs about code generated for a source text, the tool described in this paper can be seen as a very simple and limited certifying compiler [7].

7 Conclusions

An adequate encoding for sets (and by extension, maps) motivates the desire for bijections from arbitrary type families to the natural numbers. Mapping functions are defined by enumerating the instances of “finite constructors” and mapping them to small numbers. For the infinite constructors, the mapping uses the INTERCAL mingle operation, multiplication and addition to divide up the remaining numbers. HOAS variables are handled by reserving space “under” the finite constructors for variables. Proofs are generated for the four requirements of a bijection (total, unique, onto and one-to-one). Handling variables requires some Twelf-specific idioms for handling contexts. The “onto” proof requires lemmas to handle mappings to 0 and 1 to base the induction on.

Whether a mapping even exists for a given type family is undecidable, so the tool described includes a number of limitations, but it seems to be powerful enough to handle many type families used in practical syntaxes.

Acknowledgments

I thank Rob Simmons for taking up my challenge to define a provable bijection between HOAS and DeBruijn-based syntax. His work helped me understand much better how one can work with variables and contexts in Twelf. In particular, he showed how to define a case construct, with an effectiveness lemma in its own context and the requisite “fake” declarations to avoid falling afoul of auto-freezing. I thank the anonymous reviewers who gave much constructive criticism, some of which is now reflected in the paper.

SDG

References

- [1] Richard S. Bird & Ross Paterson (1999): *De Bruijn Notation as a Nested Datatype*. *Journal of Functional Programming* 9(1).
- [2] Karl Cray (2008): *Explicit Contexts in LF*. In: *3rd International Workshop on Logical Frameworks and Meta-languages: Theory and Practice*. Revised version at www.cs.cmu.edu/~cray/papers/2009/excon-rev.pdf.
- [3] Robert Harper, Furio Honsell & Gordon Plotkin (1993): *A framework for defining logics*. *Journal of the ACM* 40(1), pp. 143–184.
- [4] Robert Harper & Daniel R. Licata (2007): *Mechanizing metatheory in a logical framework*. *J. Funct. Program.* 17(4-5), pp. 613–673. Available at <http://dx.doi.org/10.1017/S0956796807006430>.
- [5] Daniel Licata & Robert Harper (2007). *An extensible theory of indexed types*. Web published. Available at <http://www.cs.cmu.edu/~rwh/papers/extidx/paper.pdf>.
- [6] The Coq development team (2008): *The Coq proof assistant reference manual*. LogiCal Project. Available at <http://coq.inria.fr>. Version 8.2.
- [7] George C. Necula & Peter Lee (1997): *The Design and Implementation of a Certifying Compiler*. In: *Proceedings of the ACM SIGPLAN '97 Conference on Programming Language Design and Implementation*, 32, ACM Press, New York, pp. 333–344.
- [8] Martin Odersky (2008). *The Scala Language Specification (Version 2.7)*. Web published. Available at <http://www.scala-lang.org/docu/files/ScalaReference.pdf>.
- [9] Frank Pfenning & Conal Elliott (1988): *Higher-order abstract syntax*. In: *Proceedings of the ACM SIGPLAN '88 Conference on Programming Language Design and Implementation*, 23, ACM Press, New York, pp. 199–208.
- [10] Benjamin C. Pierce (2002): *Types and Programming Languages*. The MIT Press, Cambridge, Massachusetts, USA and London, England.
- [11] Donald R. Woods, James M. Lyon, Louis Howell & Eric S. Raymond (1996). *The INTERCAL Programming Language: Revised Reference Manual*. Web published. Available at <http://www.catb.org/~esr/intercal/intercal.ps.gz>.

A Getting the Tool

The bijection generator is available as

<http://www.cs.uwm.edu/faculty/boyland/papers/map-natural.tar>

In order to use it, it will be necessary to first install Scala 2.7 [8] and my Twelf library (<http://www.cs.uwm.edu/faculty/boyland/proof/index.html>).